

Can anything from Noether's theorem be salvaged for discrete dynamical systems?

Silvio Capobianco * and Tommaso Toffoli

The dynamics of a physical system is linked to its phase-space geometry by Noether's theorem, which holds under standard hypotheses including continuity. Does an analogous theorem hold for discrete systems? As a testbed, we take the Ising spin model with both ferromagnetic and antiferromagnetic bonds. We show that—and why—energy not only acts as a generator of the dynamics for this family of systems, but is also conserved when the dynamics is time-invariant.

1 Introduction

In the last three hundred years, analytical mechanics has turned up a wealth of profound concepts and powerful methods for dealing with the dynamical systems of *physics*. Can these results—of which Noether's Theorem is one of the best known—be applied to the dynamics of *information systems*, which today make up such a large part of our culture, and to what extent?

Here one may anticipate two difficulties. On one hand, the above results take advantage of specific aspects of physics and may not be valid outside that territory. On the other, the methods of analytical dynamics are intimately tied to the calculus of the *continuum*, and may fail when applied to *discrete* systems such as automata and networks.

In sum, is analytical dynamics just a “bag of tricks” for doing continuum physics, or does it provide templates for tools useful in a more general context? Does the insight that it offers transcend the distinction between physical and man-made systems (cf. [12]) and that between discrete and continuous ones?

The prototypical dynamical system that we all know and love, namely, *physics*, displays a striking property. Its explicit evolution law has the form of a lookup table that to each possible *state* of the system—say, a $2n$ -tuple of real variables (or a *vector*) consisting of n positions and n momenta (where n is the number of degrees of freedom)—associates another $2n$ -tuple—its “next state.” For physical systems, it turns out that this table can be “compressed” into one that to each state associates a *single real variable*—the *energy* of that state. The first table explicitly expresses the dynamics as a *vectorial* function, while the second implicitly expresses the dynamics through a *scalar* function. This “compression” is *nondestructive*, that is, the full table can be reconstructed from the compressed table. Both the $2n$ -to-1 encoding and the converse 1-to- $2n$ encoding are achieved by simple, standardized algorithms (in analogy with “zipping” and “unzipping” a file). For this reason, the *energy function*, that is, the above energy lookup table, is said to be the *generator* of the dynamics, in the sense that plugging this function into a generic kinematic scheme (such as Hamilton's canonical equations) allows one to explicitly reconstruct a specific dynamics.

The usual way to derive the above property of energy relies on the fact that a physical system (both in classical and quantum mechanics) is imagined to evolve *continuously* in time. However, in order to have this property, being time-continuous is not enough—for this, the dynamics must have much additional structure, namely, *symplectic*, in classical mechanics, and *unitary*, in quantum.

As we generalize the concept of dynamical system, in particular to systems that, like automata of all kinds, have *discrete* state variables and evolve *discretely* in time, the question arises whether an energy-like quantity can be retained at all, and, if so, under what conditions. The first question we'll ask here is whether a cellular automaton can have something that can rightfully be called “energy,” what conditions it would have to satisfy for that, and what it would look like. We'll then ask whether and to what extent the method proposed by Noether for deriving *conserved quantities* from the *symmetries* of the dynamics of a continuous systems may be adapted to discrete systems.

2 Energy conservation

We've all heard of the principle of *conservation of energy*: In an isolated physical system, no matter what transformations may take place within it, there is a quantity, called *energy*, that remains conserved. Our belief in this principle is so strong that, whenever a violation of energy conservation is observed, our first reflex is not to question the principle, but to look more carefully into the experimental evidence and ask whether something may be escaping us. To defend the reputation of energy as a bona fide conserved quantity against new incriminating evidence, now and then physicists are even willing to stretch energy's very definition (see Feynman's peerless parable[5]).

How far shall we be able to retain a viable concept of energy as our conception of “dynamical system” keeps widening? Can energy be conserved in a *cellular automaton*? (cf. [4]) Conserved or not, should an energy be expected to be found *at all* in the latter? Is the idea of “energy” meaningful for a *Turing machine* or a *finite automaton*? In sum, what features does a dynamical system need to exhibit in order to possess a recognizable generalization of energy? Why should energy be conserved?

*This research was supported by the European Regional Development Fund (ERDF) through the Estonian Center of Excellence in Computer Science (EXCS), and by the Estonian Science Foundation under grant #7520.

3 Noether’s theorem

Nearly a century ago, Emmy Noether formulated a principle of great generality and beauty, commonly known as *Noether’s Theorem*[8, 13], which we shall state for a moment just in headline form, ignoring the fine print: “*To every symmetry of a dynamical system there corresponds a conserved quantity.*” According to this principle, the conservation of energy (a *state variable*) would be the consequence of the constancy of a system’s dynamical laws, that is, of the invariance in time of the system’s *structural parameters*.

Noether’s theorem plays a unifying role in our understanding of mechanics; according to [13], it “has become a fundamental tool of modern theoretical physics.” Since it can take advantage of many kinds of symmetries, it is very productive: for instance, it predicts that, if a system’s laws are rotationally invariant, then, just because of that structural invariance, there will be a specific state variable that is conserved, namely, *angular momentum*.

However, when we get down to reading Noether’s theorem’s fine print, we discover a number of restrictive provisos. A fuller version of the contract would spell “To every *one-parameter, continuous group* of symmetries of a *Lagrangian* dynamical system there corresponds a *scalar, real-valued* conserved quantity.” Thus, a *discrete* group of symmetries, such as the invariance of Conway’s “game of LIFE” under quarter-turn rotations of the orthogonal grid, won’t do (bye-bye angular momentum?). Similarly, even though the dynamics of LIFE is the same at every step, and thus *time-invariant*, it is not *invertible* (and thus only forms a semigroup, not a group), and in any event the game advances in discrete steps, instead of along a one-dimensional continuum (bye-bye energy?).

4 The Ising playground

The best way to identify the issues at play is to examine a specific class of models, and then map the conceptual question of our investigation to concrete features of a model. For our purposes, we shall choose a standard cellular-automaton realization of the *Ising spin model* (cf. [2, 3, 9]). This consists of a regular array of elementary magnets, or *spins*, that can assume one of two orientations—‘up’ and ‘down’.

Spins interact with one another. In the Ising model stylization of physics, long-range forces are ignored and short-range ones represented by the interaction of a spin with only its immediate neighbors: its four first neighbors, in a 2D orthogonal array. In the simplest case, these four *bonds* (or “couplings”) are all of a *ferromagnetic* nature, that is, they behave like rubber bands that are relaxed when the spins at the two ends of a band point in the same direction (“parallel”), and stretched, when in the opposite direction (“anti-parallel”).

Different kinds of dynamics may be used for the Ising model; in order to have a unifying criterion for prescribing a dynamics, one starts with assigning one unit of a notional “potential energy” to stressed bonds, and zero to unstressed ones. The class of dynamics we choose here, the so-called

microcanonical Ising model, are strictly deterministic and invertible: on a given step, a spin will *flip* (that is, reverse its orientation) *if and only if* doing so will leave the sum of the potential energies of the four surrounding bonds *unchanged*. Note that, in the ferromagnetic case, this happens when two of the neighbors are parallel to the given spin and the other two anti-parallel. In this case, in fact, we start with two stressed bonds. After the flipping of a given spin, its parallel neighbors will have turned anti-parallel and vice versa, so that the two stretched bonds relax at the same time as the two relaxed ones stretch, and consequently the overall potential energy stored in these bonds is two units both *before* and *after* the step. In all other cases, since no flip occurs, potential energy is conserved trivially.

A final stipulation. The above rule would become inconsistent if one attempted to update two adjacent spins *at the same time*. In fact, if *both* spins are instructed to flip, the assumption on the part of either spin that the shared bond’s stress status will thereupon be complemented *fails* (as its stress status will remain unchanged), and energy is no longer conserved. A standard solution is to treat the array of spins as a checkerboard consisting of two intermeshed subarrays called *even* and *odd* (the red and black squares of the board) according to the parity of $x + y$, where x and y are the spatial coordinate of a site. A system *state* s consists of an *ordered pair* $\langle q_A, q_B \rangle$ of *configurations* of different parity (i.e., one even and one odd, in either order). Together, an even and an odd configuration fill up the whole array; however, their order in the pair is relevant, that is, state $\langle q_A, q_B \rangle$ is distinct from $\langle q_B, q_A \rangle$. Depending on the context, it may be convenient to call the two configurations of a pair either the “past” and the “present” or, alternatively, the “present” and the “future” (the latter abbreviated, when subscript, as “pres” and “futr”).

The Ising dynamics is a second-order recurrence relation of the form

$$q_{x,y}^{t-1} \oplus f(q_{x-1,y}^t, q_{x+1,y}^t, q_{x,y-1}^t, q_{x,y+1}^t) = q_{x,y}^{t+1}. \quad (1)$$

This is a *conditional permutation*: a spin $q_{x,y}$ in the *past* (time $t - 1$) is mapped into a spin $q_{x,y}$ in the *future* (time $t + 1$) by XORing it with an *enable* bit—the output of the binary function f when given as arguments the *present* values (time t) of the first-neighbor spins $q_{x-1,y}, q_{x+1,y}, q_{x,y-1}, q_{x,y+1}$, which belong to the complementary subarray and thus affect the updating of $q_{x,y}$ but are not affected it. Depending on the value of the enable bit, this updating either leaves a spin unchanged or complements it—both invertible operations. The above “shifting-frame” mechanism—of the form $\langle q_{\text{past}}, q_{\text{pres}} \rangle \mapsto \langle q_{\text{pres}}, q_{\text{futr}} \rangle$, where the right item q_{pres} of the old state shifts to the left unchanged, pushing out the left item q_{past} , and a new item q_{futr} , computed according to (1), shifts in from the right—is called in numerical analysis *leapfrog updating*.

The macroscopic behavior of this kind of model is trivial in one dimension; conceptually productive in two dimensions, as it provides by very simple means an insight into the physics of phase transitions; and increasingly rich and challenging in more dimensions.

5 Why call it energy?

Setting physical motivation aside, we have in the Ising spin model an abstract symbolic dynamics consisting of a two-dimensional array over the binary alphabet $\{0, 1\}$ (for what was ‘down’ and ‘up’; it will still be convenient to call “spin” the state variable at each site). We distinguish between even and odd sites according to the parity of the sum $x+y$ of their coordinates. We update even and odd sites on alternating steps according to the rule (stated the rule in a way that obviously generalizes to any number of dimensions) that

“A spin u at $\langle x, y \rangle$ flips (i.e., complements) *if and only if* its first neighbors (those at $\langle x \pm 1, y \rangle$ and $\langle x, y \pm 1 \rangle$) are *evenly divided* between state u and its complement \bar{u} .”

Note that we had no use for an “energy.” However, if as an afterthought we choose to call *energy* a distributed quantity whose value, for each *pair* of adjacent first-neighbor spins, is 0 if these two spins are equal and 1 if different, then it follows that energy is *conserved* by the above dynamics. Also note that this energy is but the *length of the boundary* between ‘up’ and ‘down’ domains. While the *magnetization*—number of spins up minus number of spins down—may change with time, that length, and thus the energy, remains constant. Is this use of the term “energy” just physics nostalgia—like calling a British colony “New Hampshire”? Or is there something more to it?

For most people, energy is a fungible resource of the form $E = T + U$, that is, stored either in the inertia of a moving object (“kinetic energy” T) or in the state of compression of an elastic medium such as a spring (“potential energy” U). Such a concrete presentation of energy will do for many practical applications.

But if you look at what energy is *really* supposed to be in physics, you’ll discover an abstract quantity that is totally unconcerned with the *nature* of the objects and materials that make up a dynamical system, and instead totally absorbed with the topological structure of the system’s *trajectories*. In fact, the total energy of a system may be defined as

1. A *real-valued function* of the system’s *state*,
2. that is *additive*,
3. and is a *generator of the dynamics*.

Incidentally, there is no mention of “energy conservation” in all this.

NOTE. We wish we could give a precise reference for the above definition. In fact, here we rely on disparate loose hints culled and integrated with some effort from the literature. Physicists are supposed to know what energy is, and they don’t need—perhaps have good reasons not to *want*—a final definition (cf. the Feynman parable mentioned in §2). So, not only college textbooks, but even such beloved conceptually-minded references as Arnold[1], Goldstein[6], and Lanczos[7] proceed by giving *examples* of energy in increasingly elaborate contexts, and then bringing one’s attention to increasingly more abstract *properties* of energy. Definitions—hardly ever.

Let’s verify that the “energy” of the Ising model agrees with our abstract definition. In order to do so, we’ll be

forced to re-examine the definition and pore over the fine print that by tacit agreement stands behind such terms as “state,” “additive,” and “generator.”

5.1 Additivity

As we’ve seen, what we proposed to call Ising “energy” of a state is simply the *number of stretched bonds* in that state. Incidentally, since the Ising model consist of an *infinite* array of sites, the state of the system consists of an infinite array of spins, and virtually any state will have *infinite* energy. For a function to be *additive*, here the physicist means that (a) it is meaningful to partition the system into separate subsystems such that each of these subsystems has its own well-defined state, and the global state is the composition, or Cartesian product, of the states of these subsystems; (b) the function is defined on the state of each of these subsystems, and (c) its value for the whole system is the sum of its values for the individual subsystems.

If a system could be partitioned into subsystems or “blocks” that are totally independent of one another, each one having a well-defined energy independent of the state of the other blocks, additivity would be trivial to achieve—just *define* the energy of the whole system as the sum of the energy of its subsystems—but by the same token would be a *vacuous* concept. In general, however, there will be some residual coupling between blocks, and even if one knew the precise state of a block one would be able to assign an energy to it only to within an upper bound and a lower bound, which reflect the lack of knowledge of the block’s environment’s state. Fortunately, in many cases, when interactions are *local* and the blocks are appropriately chosen, this uncertainty is roughly proportional to the length of a block’s *boundary*—where it interfaces with its environment, and thus the *relative* uncertainty is bounded above by the boundary/volume ratio for a block, decreases with block size, and vanishes in the limit of arbitrarily large blocks. This is what “additivity” means to the physicist, and it is consistent with the approach taken by topological dynamics.

In the Ising model, sites interact with one another through an unbroken mesh of couplings; there are no strictly independent subsystems. However, the stipulation made at the end of §4, whereby the odd mesh is kept constant when the even mesh is updated (and vice versa)—and thus effectively acts as a parameter rather than a state-variable—enforces a discipline by which energy variations corresponding to local variations of a configuration can be determined *exactly*. We shall see this at the end of §5.2 and use it to full advantage in §7.

5.2 Generator of the dynamics

A *generator of the dynamics* is a function of the system’s state through the knowledge of which one can reconstruct the system’s dynamics in an explicit, vectorial form (cf. §1)—of course only *up to an isomorphism*. It is understood that this reconstruction should be achievable by an algorithm given once-and-for-all, independent of the specific system (otherwise, one could “cheat” and hide information

about the system in the algorithm itself—big deal!). Suppose the function is a “Hamiltonian” that for each state $\langle q, p \rangle = \langle q_1, \dots, q_n, p_1, \dots, p_n \rangle$ yields its energy $H(q, p)$. Our goal is to determine where any state $\langle q, p \rangle$ will go.

For a continuous dynamics as we imagine in the present example one would want the instantaneous “direction and rate” at which the state progresses, that is, the $2n$ derivatives $\langle \dot{q} = \frac{dq}{dt}, \dot{p} = \frac{dp}{dt} \rangle$. In other words, we want to arrive at *vectorial* mapping, of the form

$$\langle q, p \rangle \mapsto \langle \dot{q}, \dot{p} \rangle, \quad (3)$$

even though mere evaluation of the function at $\langle q, p \rangle$ will perform the mapping $\langle q, p \rangle \mapsto H(q, p)$ and thus only give us a scalar. However, if we performed *repeated* samplings of H in the vicinity of $\langle q, p \rangle$ (after all, we have the entire function H and we can apply it as many time as we want) so as to get a “sense of direction,” or, equivalently, in the continuum case, if we could get the *derivatives* of H at that point, we would get more information associated with that point than its mere energy $H(q, p)$. In Hamiltonian mechanics, mapping (3) is provided by the standard *canonical equations*

$$\dot{q} = \frac{\partial H}{\partial p}; \quad \dot{p} = -\frac{\partial H}{\partial q}. \quad (4)$$

Note that, since q and p are actually abbreviations for n -tuples, here we have $2n$ equations yielding a *vector* consisting of n pairs $\langle \dot{q}_i, \dot{p}_i \rangle$, ($i = 1, \dots, n$).

It will presently become obvious that in the Ising model the energy function *does* generate the dynamics—and by means of a very simple recipe. In a *discrete* Hamiltonian dynamics, a state is no longer a “position/momentum” pair $\langle q, p \rangle$ as in the continuous case, but an ordered *pair of configurations* $\langle q_0, q_1 \rangle$, as explained at the end of §4. The rules of the game are that one may propose *any* such state, as *many times as one wishes*, and each time get as an answer the energy $H(q_{\text{past}}, q_{\text{pres}})$ for that state. Isn’t this number, in the range $\{0, 1, 2, \dots, \infty\}$ (and almost always ∞) much too little to determine the infinite q_{futr} configuration? The answer lies in the principle of *virtual displacements*. That is, given q_{past} and q_{pres} , one may propose *any conceivable candidate* for q_{futr} , ask for the energy of the resulting new state $\langle q_{\text{pres}}, q_{\text{futr}} \rangle$, compare it with the old energy, and on the basis of this comparison decide whether to accept or reject the candidate. Just as in the “hot, hot, cold, cold” game, in the game of “battleships”—or, for that matter, in biological evolution—one is never explicitly shown the target. Instead, one must describe an object, and will only be told (“yes,” “no”) whether that is target, or possibly given an intermediate scalar value (“hot”, “warm”, etc.) that indicates *how far* one’s proposal is from the target. In principle, one can arrive at the solution by going through all possible candidates. The point of the game is to speed up the search by dynamically arranging the proposals so that the feedback will give one a “sense of direction.”

In the Ising case, we want to find a q_{futr} such that the energy difference $\Delta E = H(q_{\text{pres}}, q_{\text{futr}}) - H(q_{\text{past}}, q_{\text{pres}})$ is zero not only globally but also *locally*, that is, at any scale down to the radius of the first-neighbor neighborhood (the range

of interspin “interactions”). Forget about getting finite energies by taking finite configurations by blocking (§5.1); just take the 32 *infinite* configurations of the form

$$\begin{array}{cccccc} \vdots & \vdots & \vdots & \vdots & \vdots & \\ \dots & 0 & 0 & 0 & 0 & 0 & \dots \\ \dots & 0 & 0 & n & 0 & 0 & \dots \\ \dots & 0 & w & c & e & 0 & \dots \\ \dots & 0 & 0 & s & 0 & 0 & \dots \\ \dots & 0 & 0 & 0 & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \end{array} \quad (5)$$

with c denoting a “center cell” in the past and n, s, w, e denoting its four neighbors in the present. This infinite state will have a finite energy ranging from 0 through 4. Let’s make a table of energy for the 32 possible collective values of those five sites,

$cnswe$	$H f c'$	$cnswe$	$H f c'$
0 0000	0 0 0	1 0000	4 0 1
0 0001	1 0 0	1 0001	3 0 1
0 0010	1 0 0	1 0010	3 0 1
0 0011	2 1 1	1 0011	2 1 0
0 0100	1 0 0	1 0100	3 0 1
0 0101	2 1 1	1 0101	2 1 0
0 0110	2 1 1	1 0110	2 1 0
0 0111	3 0 0	1 0111	1 0 1
0 1000	1 0 0	1 1000	3 0 1
0 1001	2 1 1	1 1001	2 1 0
0 1010	2 1 1	1 1010	2 1 0
0 1011	3 0 0	1 1011	1 0 1
0 1100	2 1 1	1 1100	2 1 0
0 1101	3 0 0	1 1101	1 0 1
0 1110	3 0 0	1 1110	1 0 1
0 1111	4 0 0	1 1111	0 0 1

where H is the “Hamiltonian energy” of the Ising model for the entire global state (5), f the “conditioning” function that appears in (1), and c' , as we shall see in a moment, the effective future value of the center site. As we recall from §4, the Ising dynamics is “flip if and only if that leaves the energy of the neighborhood unchanged.” Therefore, for any state of the form (5), we may propose as candidates for the new state either (a) the same state, i.e., with $c' = c$, or (b) the one where c has flipped, i.e., where $c' = \bar{c}$. By using table (6) *two* times, i.e., for the two c' candidates c and \bar{c} , we see that the energy is (obviously) unchanged in the first case, and *may* have changed in the second case. If the energy is the *same* in both cases, then the *actual* value for c' (that is, the one prescribed by the Ising dynamics) will be the default value c ; otherwise, it will be \bar{c} . This is the value given in column c' of (6).

Finally, we must show that this new value c' , given the old value of the entire neighborhood w_{se}^n , is independent of all the others sites of the global state—those assigned 0 in global state(5). In fact, by the definition, in the Ising model, of energy as a *sum over bonds* (or adjacent-spin pairs), the energy spanning, say, the “extended patch” w_{se}^{nar} is the sum of the energies spanning the two elementary patches w_{se}^n and n_{e}^u ; in other words, energy is *strictly additive*, and the *local dynamics* table (6) can be applied separately to every elementary patch.

Thus, the “energy” of the Ising model satisfies the abstract, “model-independent” definition of *energy* given in

(2). ... Or *does* it?

6 Raising the kinematics bar

Indeed the Ising energy (a) is a real-valued function of the system’s state; (b) it is additive, since the alternating updating scheme guarantees that the accounting of energy can be done independently on every elementary patch; and (c) a simple rule allows us to reconstruct the dynamics by sampling the energy function in the (topological) neighborhood of a state.

However, point (c) is in this case *vacuous*, since here we are dealing with a *single* system—the Ising model—and there is no indication of how one would go about generating the dynamics of *other* systems using other energy functions. Consider a telegraph channel with a single input key, pressing which will invariably transmit the text of the *Kamasutra*. In spite of the length and the intrinsic interest of this text, the capacity of the channel is obviously $\log 1 = 0$ bits. In a similar way, the information contents of H_{ising} as a function, for the sake of specifying a dynamics, is zero, because it is a single “key”: the Ising dynamics is the only dynamics we have in mind, it is known, and we might as well have stored a copy of it at the channel’s receiving end. What one needs is a generic function scheme—a look-up table “blank form,” as it were—which, when filled with specific contents and processed by an algorithm given once and for all, will allow one to describe any one of an agreed upon class of dynamical systems.

This issue is addressed by physicists by distinguishing between “dynamics” and “kinematics,” where the latter term is used in a somewhat specialized sense. That is, a *kinematics* is a framework—a set of rules—that characterizes a whole class of dynamical systems of a similar type or “texture,” while a *dynamics* is any one of the systems sharing that kinematics. Until we place the Ising model within a larger class of models, and show the existence of an energy-function-valued *variable*—whose values will be specific energy functions each of which will in turn allow a fixed algorithm to generate the different dynamics of that class—until then we will not have a bona fide *generator of the dynamics*.

There is a natural and a very simple generalization of the Ising model (used, specifically, in the modeling of *spin glasses*) that yields a class of *sixteen* dynamics. We’ll show that for that class the energy function plays a nonvacuous role as a generator of the dynamics.

In §4, we called *ferromagnetic* a bond between adjacent spins that is unstressed when the two spins are parallel and stressed when antiparallel (“type 0” bond). We now introduce a bond of “type 1,” or *antiferromagnetic*, which is stressed when the spins are *parallel*, and unstressed otherwise. With bonds of two kinds thus available, and four bonds surrounding a spin, there are $2^4 = 16$ possible *structural environments* by which a spin can find itself surrounded. We shall retain our definition of *Ising energy* as the sum of the energy of the bonds, recalling only that the stress $u \xrightarrow{b} v$ of a bond of type b between a pair of parallel spins u and v will $u \oplus b \oplus v$, rather than just $u \oplus v$ as in the

original Ising model. As part of our kinematics—not as a peculiarity of an individual dynamics—we shall retain the Ising prescription that a spin will flip if and only if such a move leaves the energy unchanged for a given assignment of the four surrounding bond types. However, within this kinematics, we shall for the moment restrict our attention to systems whose structural parameters (those determining a dynamics, and in our case consisting of the assignment of bond types) be *time- and space-invariant*, that is, for each specific dynamics of our kinematics, the behavior of a spin shall not depend on *where* the spin is located and on *what time it is*, but only on the values of its neighbors.

With the latter restriction, we are left with as many *global* assignments of structural parameters as there are choices of *local* assignments—that is, the 16 distinct assignments of four bond types around a spin. For each of these structures we must tabulate, by running over all possible states, the corresponding dynamics (a function from old global state to new global state) and the corresponding energy function (a function from global state to an integer). Observe that two distinct structures may well end up yielding the same dynamics or the same energy function. Our goal will be achieved if (a) no two dynamics are specified by the *same* energy function (in other words, the energy function *distinguishes* between dynamics), and (b) every dynamics has *at least one* energy function that specifies it. Since, by construction, energy is additive down to the scale of individual four-bond “patches,” the local “variational principle” employed in §5.2 will then allow us not only to indirectly denote, but to explicitly *generate* the corresponding vectorial-form (that is, state-to-state) dynamics.

This is indeed the case. Since we are presenting here an original approach, it is important not to deceive ourselves on delicate aspects of an argument. For this reason, as a “sanity check” we have explicitly tabulated dynamics and energy function for all 16 elements of our generalized Ising kinematics on a finite, toroidally wrapped-around model. (We used just a 4×4 torus; though small, this size is large enough to rule out spurious degeneracies. Full scripts, output data, and statistics are available at www.ioc.ee/~silvio/nrg/.) A formal proof is subsumed within the properties of the numerical integration scheme of Fig. 1, which constitute a proof in the more general context of space- and time-dependent dynamics. In fact, in the next section we shall remove the constraint that the structure—and the attendant dynamics—be spacetime uniform.

7 Space- and time-dependent dynamics

Energy in these generalized Ising models is a function of the global state—a function ultimately specified by the structure parameters of the specific model. If the parameters are allowed to change from one step to the next, then this function is in turn a function of *time*. In fact, the type of the bond joining two given spins will depend on “what time it is” at the moment of performing a transition $\langle q_A, q_B \rangle \rightarrow \langle q_B, q_C \rangle$. The “new state” $\langle q_B, q_C \rangle$ at the end of this transition reappears as the “old state” at the

beginning of the following transition, $\langle q_B, q_C \rangle \rightarrow \langle q_C, q_D \rangle$, but now it is a *different time*, and the energy of the *same state* (of spins) has to be re-evaluated accordingly to the new assignment of structural parameters.

This arrangement of things, which is essentially a numerical integration scheme, is graphically illustrated by Fig. 1. There, the integers along the time axis represent discrete time instants, those along x or y , discrete spatial positions. A spin is represented by a thick vertical line, interrupted by a conditional-permutation gate (\oplus) every time that it is the spin's turn to be updated. A system state is an ordered pair of configurations associated with two consecutive integer times. This state may conveniently be labeled with the “half-time” in between; that is, one may denote by $\frac{1}{2}$ the state consisting of the configurations at 0 and 1, and so forth. Steps of the dynamics are nominally associated with integer values of time. So, for instance, state $1\frac{1}{2}$ is the “new state” for step 1 but the “old state” for step 2. Horizontal arcs of unit length (thin lines) represent interspin bonds as at the time of a step; letters may be used to specify the nature of a bond; so r denotes the nature of the bond between spins 1 and 2 at step 1, while s denotes the nature of the same bond at step 2.




Figure 1: (a) Spacetime spin grid (y not shown). (b) Detail of full grid, with both x and y .

The value of a bond remains steady during spin updating (at integer times) but may change between updatings (at half-integer times). On this occasion, the energy of a state is accordingly re-evaluated. Thus, even if energy is conserved during spin updating, in a time-dependent system the energy may still change when the *bonds* are altered.

Thus, it will generally happen that energy changes with time. With the present *full* class of dynamics (i.e., one which allows arbitrary spacetime variation of the structural parameters) of the given generalized Ising kinematics, we arrive at the conclusion that energy is *necessarily conserved if* the spacetime structure is time-invariant (because the transitions themselves are by construction energy conserving, and with time-invariant bond types no re-evaluation of the energy need occur). On the other hand, when the spacetime structure is *not* time invariant, *at least some* of the dynamics are not energy-conserving. We can thus rightfully say that, in the present context, “Ising energy” as a *whole* (not the energy of a specific Ising dynamics) is the quantity that is conserved *because of* the time-invariance of the dynamics.

For this class of dynamics, then, in spite of the discreteness of the dynamics, a non-frivolous transliteration of Noether’s theorem applies with full force, yielding energy conservation from time invariance.

8 Conclusions

We’ve shown that certain aspects of Noether’s theorem apply to dynamical systems beyond ordinary analytical mechanics, and *why they do*. We were able to shed the requirement that the dynamics be *continuous*; still, some form of *second-order* discipline seems to remain essential for a symmetry to give rise to a conservation law. This is just the beginning of what promises to be a productive line of research.

List of references

- [1] ARNOLD, Vladimir, *Mathematical Methods of Classical Mechanics*, 2nd ed., corr. 4th printing, Springer 2010.
- [2] BACH, Ted, “Methodology and implementation of a software architecture for cellular and lattice-gas automata programming,” PhD thesis, Boston University 2007; pm1.bu.edu/~tt/Bach07thesis.pdf.
- [3] BENNETT, Charles H, Norman MARGOLUS, and Tommaso TOFFOLI, “Bond-energy variables for Ising spin-glass dynamics,” *Phys. Rev. B* **37** (1988), 2254.
- [4] BOYKETT, Tim, Jarkko KARI, and Siamak TAATI, “Conservation Laws in Rectangular CA,” *J. Cell Autom.* **3** (2008), 115–122.
- [5] FEYNMAN, Richard, Robert LEIGHTON, and Matthew SANDS, “Conservation of Energy,” *The Feynman Lectures on Physics*, Vol. 1, Addison-Wellesly 1963, Ch. 4 (4-1–4-8).
- [6] GOLDSTEIN, Herbert, Charles POOLE, and John SAFKO, *Classical Mechanics*, 3rd ed., Addison-Wesley 2001.
- [7] LANCZOS, Cornelius, *The Variational Principles of Mechanics*, 4th ed., Dover 1986.
- [8] NOETHER, Emmy, “Invariante Variationsprobleme,” *Nachr. D. König. Gesellsch. D. Wiss. Zu Göttingen, Math-phys. Klasse* 1918, 235–257 (English translation at arxiv.org/abs/physics/0503066v1).
- [9] POMEAU, Yves, “Invariant in cellular automata,” *J. Phys. A* **17** (1984), L415–L418.
- [10] TOFFOLI, Tommaso, “Action, or the fungibility of computation,” in *Feynman and Computation* (Anthony HEY, ed.), Perseus (1999), 349–392.
- [11] TOFFOLI, Tommaso, Silvio CAPOBIANCO, and Patrizia MENTRASTI, “A new inversion scheme, or how to turn second-order cellular automata into lattice gases,” *Theor. Comp. Sci.* **325** (2004), 329–344.
- [12] TYAGI, Akhilesh, “A principle of least computational action,” *Workshop on Physics and Computation*, IEEE Computer Society Press 1994, 262–266.
- [13] WIKIPEDIA, “Noether’s theorem,” [en.wikipedia.org/wiki/Noether’s_theorem](http://en.wikipedia.org/wiki/Noether's_theorem) (sampled on 19 Jan. 2011).